

Fitting Power-laws in empirical data with estimators that work for all exponents

Rudolf Hanel¹, Bernat Corominas-Murtra¹, Bo Liu¹, and Stefan Thurner^{1,2,3,4*}

¹*Section for Science of Complex Systems, Medical University of Vienna, Spitalgasse 23, 1090 Vienna, Austria*

²*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA*

³*IIASA, Schlossplatz 1, 2361 Laxenburg, Austria and*

⁴*Complexity Science Hub Vienna, Josefstädterstrasse 39, A-1090 Vienna, Austria*

(Dated: September 20, 2016)

It has been repeatedly stated that maximum likelihood (ML) estimates of exponents of power-law distributions can only be reliably obtained for exponents smaller than minus one. The main argument that power laws are otherwise not normalizable, depends on the underlying sample space the data is drawn from, and is true only for sample spaces that are unbounded from above. Here we show that power-laws obtained from bounded sample spaces (as is the case for practically all data related problems) are always free of such limitations and maximum likelihood estimates can be obtained for arbitrary powers without restrictions. Here we first derive the appropriate ML estimator for arbitrary exponents of power-law distributions on bounded discrete sample spaces. We then show that an almost identical estimator also works perfectly for continuous data. We implemented this ML estimator and discuss its performance with previous attempts. We present a general recipe of how to use these estimators and present the associated computer codes.

INTRODUCTION

The omnipresence of power-laws in natural, socio-economic, technical, and living systems has triggered immense research activity to understand their origins. It has become clear in the past decades that there exist several distinct ways to generate power-laws (or asymptotic power-laws), for an overview see for example [1, 2]. In short, power-laws of the form

$$p(x) = Cx^{-\lambda}, \quad (1)$$

arise in critical phenomena [3, 4], in systems displaying self-organized criticality [5], preferential attachment type of processes [6–9], multiplicative processes with constraints [10], systems described by generalized entropies [11, 12], or sample space reducing processes [13], i.e. processes that reduce the number of possible outcomes (sample space) as they unfold. Literally thousands of physical, natural, man-made, social, and cultural processes exhibit power-laws, the most famous being earthquake magnitudes [14, 15], city sizes [16, 17], foraging and distribution pattern of various animal species [18], evolutionary extinction events [19], or the frequency of word occurrences in languages, known as Zipf’s law [20].

It is obvious that estimating power-law exponents from data is a task that sometimes should be done with high precision. For example if one wants to determine the universality class a given process belongs to, or when one estimates probabilities of extreme events. In such situations small errors in the estimation of exponents may lead to dramatically wrong predictions with potentially serious consequences.

Estimating power-law exponents from data is not an entirely trivial task. Many reported power-laws are simply not exact power-laws, but follow other distribution functions. Despite the importance of developing adequate methods for distinguishing real power-laws from

alternative hypotheses, we will not address this issue here since good standard literature on the topic of Bayesian *alternative hypotheses testing* exists, see for example [21, 22]. For power-laws some of these matters have been discussed also in [24]. Here we simply focus on estimating power-law exponents from data on a sound probabilistic basis, using a classic Bayesian parameter estimation approach, see e.g. [23], that provides us with *maximum likelihood* (ML) estimators for estimating power-law exponents over the full range of reasonably accessible values. Having such estimators is of particular interest for a large classes of situations where exponents close to $\lambda \sim 1$ appear (Zipf’s law). We will argue here that whenever dealing with data we can assume discrete and bounded samples spaces (domains), which guarantees that power-laws are normalizable for arbitrary powers λ . We then show that the corresponding ML estimator can then also be used to estimate exponents from data that is sampled from continuous sample spaces, or from sample spaces that are not bounded from above.

Questions before fitting power-laws

In physics the theoretical understanding of a process sometimes provides us with the luxury of knowing the exact form of the distribution function that one has to fit to the data. For instance think of critical phenomena such as Ising magnets in 2 dimensions at the critical temperature, where it is understood that the susceptibility follows a power-law of the form $(T - T_c)^{-\gamma}$, with γ a critical exponent, that occasionally even can be predicted mathematically. However, often – and especially when dealing with complex systems – we do not enjoy this luxury and usually do not know the exact functions to fit to the data.

In such a case, let us imagine that you have a data

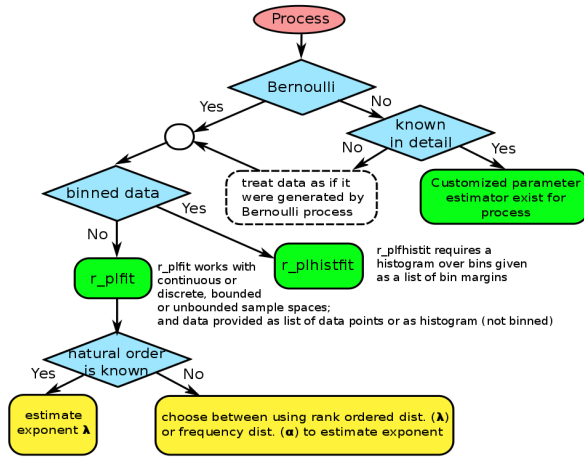


FIG. 1. Decision tree of questions that should be clarified before estimating power-law exponents from data. The tree shows under which conditions the fitting algorithms developed in this paper `r_plfit` and `r_plhistfit` can be used.

set and from first inspection you think that a power-law fit could be a reasonable thing to do. It is then essential, before starting with the fitting procedures, to clarify what one knows about the process that generated this data. The following questions may help to do so.

- Do you have information about the dynamics of the process that is generating what appears to be a power-law?
- Is the data generated by a Bernoulli process (e.g. tossing dice), or not (e.g. preferential attachment)?
- Is the data available as a collection of samples (a list of measurements), or only coarse-grained in form of a histogram (binned or aggregated data).
- Is the data sampled from a discrete (e.g. text) or continuous sample space (e.g. earthquakes)?
- Does the data have a natural ordering (e.g. magnitudes of earthquakes), or not (e.g. word frequencies in texts)?

The decisions one has to take before starting to estimate power-law exponents are shown as a decision-tree in Fig. (1). If it is known that the process generating the data is not a Bernoulli process (for example if the process belongs to the family of history dependent processes such as e.g. preferential attachment), then one has the chance to use this information for deriving parameter estimators that are tailored exactly for the particular family of processes. If no such detailed information is available one can only treat the process as if it were a Bernoulli process, i.e. information about correlations between samples is ignored. If we know (or assume) that the data generation process is a Bernoulli process, the next thing to

determine is whether the data is available as a collection of data points, or merely as coarse grained information in form of a histogram that collects distinct events into bins (e.g. histograms of logarithmically binned data).

If data is available in form of a data set of samples (not binned), a surprisingly general maximum likelihood (ML) estimator can be used to predict the exponent of an underlying power-law $p(x) \propto x^{-\lambda}$. This estimator that we refer to as ML^* , will be derived in the main section. Its estimates for the underlying exponent λ , are denoted by λ^* . The code for the corresponding algorithm we refer to as `r_plfit`. If information is available in form of a histogram of binned data, a different estimator becomes necessary. The corresponding algorithm (`r_plhistfit`) is discussed in APPENDIX A and in the section below on discrete and continuous sample spaces. Both algorithms are available as matlab code [32]. For how to use these algorithms, see APPENDIX B.

If we have a dataset of samples (not binned), so that the `r_plfit` algorithm can be used, it still has to be clarified whether the data has a natural order or not? Numerical observables such as earthquake magnitudes are *naturally* ordered. One earthquake is always stronger or smaller than the other. If observables are non-numeric, such as word types in a text, then a natural order can not be known *a priori*. The natural order can only be inferred approximately by using so-called *rank-ordering*; or alternatively – by using the so-called *frequency distribution* of the data. Details are discussed below in the section on rank-order, frequency distributions, and natural order.

Other issues to clarify are to see if a given sample space is continuous or discrete, and if the sample space is bounded or unbounded. These questions however, turn out to be not critical. One might immediately argue that for unbounded power-law distribution functions normalization becomes an issue for exponents $\lambda \leq 1$. However, this is only true for Bernoulli processes on *unbounded* sample spaces. Since all real-world data sets are collections of finite discrete values one never has to actually deal with normalization problems. Moreover, since most experiments are performed with apparatus with finite resolution, most data can be treated as being sampled from a bounded, discrete sample space, or as binned data. For truly continuous processes the probability of two sampled values being identical is zero. Therefore, data sampled from continuous distributions can be recognized by sample values that are unique in a data set. See APPENDIX A for more details.

Statistically sound ways to fit power-laws were advocated and discussed in [24–26, 31]. They overcome intrinsic limitations of the *least square* (LS) fits to logarithmically scaled data, which were and are widely (and often naively) used for estimating exponents. The ML estimator that was presented in [24] we refer to as the ML_{CSN} (for Clauset-Shalizi-Newman) estimator; its estimates for the exponent we denote by $\hat{\lambda}$. The approach

that leads to ML_{CSN} focuses on continuous data x that follows a power-law distribution from Eq. (1), and that is bounded from below $x > x_{\min} > 0$ but is not bounded from above (i.e. $x_{\max} > x$ with $x_{\max} = \infty$). In [24] emphasis is put on how ML estimators can be used to infer whether an observed distribution function is likely to be a power-law or not. Also the pros and cons of using cumulative distribution functions for ML estimates are discussed, together with ways of treating discrete data as continuous data. For the continuous and unbounded case, simple explicit equations for the ML_{CSN} estimator can be derived. The continuous approach however, even though it seemingly simplifies computations, introduces unnecessary self-imposed limitations with respect to the range of exponents that can be reliably estimated. ML_{CSN} works brilliantly for a range of exponents between -3.5 and -1.5 .

Here we show how to overcome these limitations – and by doing so extend the accessible range of exponents – by presenting the exact methodology for estimating λ for discrete bounded data with the estimator ML^* . While this approach appears to be more constrained than the continuous one we can show also theoretically that data from continuous and potentially unbounded sample spaces can be handled within essentially the same general ML framework as well. The key to the ML^* estimator is that it is not necessary to derive explicit equations for finding λ^* . Implicit equations in λ exist for power-law probability distributions over discrete or continuous sample spaces that are both bounded from below and above. Solutions λ^* can be easily obtained numerically. An implementation of the respective algorithms can be found in [32], for a tutorial see APPENDIX B.

Rank-order, frequency distributions & natural order

There exist three distinct types of distribution functions that are of interest in the context of estimating power-law exponents:

- i: The *probability distribution* $p(x)$ assigns a probability to every observable state-value x . Discrete and bounded sample spaces are characterized by W state-types $i = 1, \dots, W$, with each type i being associated with a distinct value $x = z_i$.
- ii: The *relative frequencies*, $f_i = k_i/N$, where k_i is the number of times that state-type i is observed in N experiments. $k = (k_1, \dots, k_W)$ is the *histogram* of the data. As explained below in detail, the relative frequencies can be ordered in two ways.
 - If f_i is ordered according to their descending magnitude this is called the *rank ordered* distribution.

- If f_i is ordered according to the descending magnitude of the probability distribution $p(z_i)$, then they are *naturally ordered* relative frequencies.

- iii: The *frequency distribution* $\phi(n)$ counts how many state-types i fulfill the condition $k_i = n$.

In Fig. (2) we show these distribution functions. There $N = 10000$ data points are sampled from $x \in \{1, \dots, 1000\}$, with probabilities $p(x) \propto x^{-0.7}$. The probability distribution is shown (red). The relative frequency distribution f is plotted in natural order (blue), the rank-ordered distribution is shown with the yellow line, which clearly exhibits an exponential decay towards the tail. The inset shows the frequency distribution $\phi(n)$ of the same data. We next discuss how different sampling processes can be characterized in terms of natural order, rank-order, or frequency distributions.

Processes with naturally ordered observables

For some sampling processes the ordering of the observed states is known. For example think of x representing the numerical values of earthquake magnitudes. Here any two observations x and x' can be ordered with respect to their numerical value, or their *natural order*. Since power-law distributions $p(x) \propto x^{-\lambda}$ are monotonic this is equivalent to ranking observations according to the probability distribution p they are sampled from: The most likely event has *natural rank* 1, the second most likely rank 2, etc. In other words, we can order state-types x in a way that over the sample space $\Omega = \{z_i | i = 1, \dots, W\}$, $p = (p(z_1), \dots, p(z_W))$ is a monotonic and decreasing function.

Processes with rank-ordered observables

If p is not known *a priori* because the state-types i have no numerical values z_i attached, as happens for example with words in a text, we can only count relative frequencies f_i (a normalized histogram) of states of type i , *a posteriori*, i.e. after sampling. To be clear, let $k = (k_1, \dots, k_W)$ be the histogram of N recorded states. k_i is the number of times we observed type i , then $f_i = k_i/N$ is the relative frequency of observing states of type i . After all samples are taken, one can now order states with respect to f_i , such that the rank 1 is assigned to state i with the largest f_i , rank 2 to i' with the second largest $f_{i'}$, etc. $f = (f_1, \dots, f_W)$ is called the *rank-ordered* distribution of the data.

The natural order imposed by p and the rank-order imposed by f are not identical for finite N . However, if data points have been sampled independently, then f converges toward p (for $N \rightarrow \infty$) and the rank-order induced by f will asymptotically approach the natural




FIG. 2. The four types of distribution functions. Data is sampled from a power-law distribution $p(x) \propto x^{-\lambda}$ with an exponent $\lambda = 0.7$ (red line). The relative frequencies f_i are shown for $N = 10000$ sampled data points according to their natural (prior) ordering that is associated with p (blue). The rank-ordered distribution (posterior) is shown in yellow, where states i are ordered according to their observed relative frequencies f_i . The rank-ordered distribution follows a power-law, except for the exponential decay that starts at rank ~ 500 . A low frequency cut-off should be used to remove this part for estimating exponents. The inset shows the frequency distribution $\phi(n)$ that describes how many states x appear n times (green). The frequency distribution has a maximum and a power-law tail with exponent $\alpha = 1 + 1/\lambda \sim 2.43$. To estimate α , one should only consider the tail of the frequency distribution function.

order induced by p . The highest uncertainty on estimating the order induced by p using f is associated with the least frequent observations. Therefore, when estimating exponents from rank-ordered distributions, one might consider to use a low-frequency cut-off to exclude infrequent data.

Frequency distributions

Exponents of power-laws can also be estimated from frequency distributions $\phi(n)$. These counts how many distinct state-types i occur exactly n times in the data. It does not depend on the natural (prior) order of states and therefore is sometimes preferred to the (posterior) rank-ordered distribution. However, complications may appear also when using $\phi(n)$. The frequency distribution $\phi(n)$ that is associated with a power-like probability distribution $p \propto x^{-\lambda}$ (and asymptotically to f) is not an exact power-law but a non-monotonic distribution (with a maximum). Only its tail decays as a power-law, $\phi(n) \propto n^{-\alpha}$. The exponents λ and α are related through the well known equation

$$\alpha = 1 + 1/\lambda \quad . \quad (2)$$

If the probability distribution has exponent λ , the tail of the associated frequency distribution has exponent α . Since the frequency distribution behaves like a power-law only in its tail, estimating α makes it necessary to constrain the observed data to large values of n . Note that this is equivalent to using a low-frequency cut-off. One option to do that is to derive a maximum entropy functional for $\phi(n)$ and fit the resulting (approximate) max-ent solution to the data. We do not follow this route here.

If the natural order of the data is known, one can directly use the natural ordered data in the ML estimates for the exponents. If it is not known, either the rank-ordered distribution can be used to estimate λ , or the frequency distribution to estimate α , see Fig. (1).

One might also estimate both, λ in the rank ordered distribution, and α in the frequency distribution of the data. Using Eq. (2) to compare the two estimates may be used as a rough quality-check. If estimates do not reasonably coincide one should check whether the used data ranges have been appropriately chosen. If large discrepancies remain between α and $1 + 1/\lambda$ this might indicate that the observed distribution function in question is only an approximate power-law, for which Eq. (2) need not hold. For a tutorial on how to use `r_plfit` to perform estimates see APPENDIX B.

Discrete and continuous sample spaces & normalization

Data can originate from continuous sample spaces $\Omega_c = [x_{\min}, x_{\max}]$, or discrete ones $\Omega_d = \{z_1, z_2, \dots, z_W\}$. To each state-type $i = 1, \dots, W$, there is assigned a state-value z_i . Whether a distribution function $p(x) = Z_\lambda^{-1} x^{-\lambda}$, with $x \in \Omega$, is normalizable or not, can only be decided once the sample space Ω has been specified. The normalization factors for continuous and discrete Ω are

$$\begin{aligned} Z_\lambda(\Omega_c) &\equiv \int_{x_{\min}}^{x_{\max}} dx x^{-\lambda} = \frac{x_{\max}^{1-\lambda} - x_{\min}^{1-\lambda}}{1-\lambda} \\ Z_\lambda(\Omega_d) &\equiv \sum_{x \in \Omega_d} x^{-\lambda} = \sum_{i=1}^W z_i^{-\lambda} \quad . \end{aligned} \quad (3)$$

For bounded sample spaces with $0 < x_{\min} \leq x \leq x_{\max} < \infty$, power-laws are always normalizable for arbitrary exponents λ , and a well defined ML estimator of λ^* exists (see below). The normalization constants in Eq. (3) can be specified in `r_plfit` (see APPENDIX B).

Data sampled from a continuous sample space Ω_c can essentially be treated as if it were sampled from a discrete sample space Ω_d , where $x \in \Omega_d$ are given by the unique collection of distinct values in the data set. That is, the data set $x = (x_1, \dots, x_N)$ contains N data points x_n (that have W unique values z_i , the states of type i) which we collect in the discrete sample space $\Omega = \{z_1, \dots, z_W\}$.

For truly continuous data we have $N = W$, since the probability of $x_m = x_n$ for $n \neq m$ is vanishing. As a consequence the histogram k_i , which counts the number of times z_i appears in the data, is essentially given by $k_i = 1$ for all $i = 1, \dots, W$. This provides us with a practical criterion for when to use the normalization constant for discrete or continuous data. For details see APPENDIX A.

The equation for the ML estimator ML^* , that yields the estimate λ^* , only requires the knowledge of the relative frequency distribution $f_i = k_i/N$ (in natural- or rank-order) of the observed state-types i , as we will see in Eq. (9) below. Therefore `r_plfit` can work either with data sets x or histograms k over the unique values in the data sets. If data comes in coarse grained form, i.e. histograms, where each bin may contain a whole range of observable values x , then an estimator is required that is different from ML^* [25], see also APPENDIX A. The corresponding code `r_plhistfit` can also be downloaded from [32].

THE ML^* -ESTIMATOR FOR POWER-LAWS FROM DISCRETE SAMPLE SPACES

Consider a family of random processes Y that is characterized by the parameters $\theta = (\theta_1, \dots, \theta_R)$. Let Y be defined on a discrete sample space $\Omega = \{z_1, z_2, \dots, z_W\}$, with $0 < z_i < \infty$. The process Y samples values $x \in \Omega$ with probability,

$$p(x|\theta, \Omega) \quad . \quad (4)$$

Let us repeat the process Y in N independent experiments to obtain a data set $y = (y_1, \dots, y_N)$. $k = (k_1, \dots, k_W)$ is the histogram of the events recorded in y , i.e. k_i is the number of times z_i appears in y . Note that $\sum_{i=1}^W k_i = N$. As a consequence of independent sampling, the probability to sample exactly k is,

$$P(k|\theta, \Omega) = \binom{N}{k} \prod_{i=1}^W p(z_i|\theta, \Omega)^{k_i} \quad , \quad (5)$$

where $\binom{N}{k} = N! / \prod_{i=1}^W k_i!$ is the multinomial factor. Bayes' formula allows us to get an estimator for the parameters θ ,

$$P(\theta|k, \Omega) = P(k|\theta, \Omega) \frac{P(\theta|\Omega)}{P(k|\Omega)} \quad . \quad (6)$$

Obviously, $P(k|\Omega) = \int d\theta P(k|\theta, \Omega) P(\theta|\Omega)$ does not depend on θ . Without further available information we must assume that the parameters θ are uniformly distributed between their upper and lower limits. As a consequence, $P(\theta|\Omega)$ also does not depend on θ within the limits of the parameter range and can be treated as a constant[?]. From Eq. (6) it follows that the value θ^* that maximizes $P(\theta|k, \Omega)$ also maximizes $P(k|\theta, \Omega)$. The

most likely parameter values $\theta^* = (\theta_1, \dots, \theta_R)^*$ are now found by maximizing the log-likelihood,

$$\begin{aligned} 0 &= \frac{\partial}{\partial \theta} \frac{1}{N} \log P(\theta|k, \Omega) \\ &= \sum_{i=1}^W f_i \frac{\partial}{\partial \theta_r} \log p(z_i|\theta, \Omega) \\ &= - \frac{\partial}{\partial \theta_r} H_{\text{cross}}(f||p(z|\theta, \Omega)) \quad , \end{aligned} \quad (7)$$

for all parameters $r = 1, \dots, R$. Here $H_{\text{cross}}(f||p(z|\theta, \Omega)) \equiv - \sum_{i=1}^W f_i \log p(z_i|\lambda, \Omega)$, is the so-called *cross-entropy*. In other words, ML -estimates maximize the cross-entropy with respect to the parameters θ_r .

The ML^* -algorithm for power-laws

To apply Eq. (7) for ML -estimates of power-law exponents, one specifies the finite sample space $\Omega = \{z_1, z_2, \dots, z_W\}$, and the family of probability density functions is,

$$p(x|\lambda, \Omega) = \frac{x^{-\lambda}}{Z_\lambda(\Omega)} \quad , \quad (8)$$

with $x \in \Omega$. Note that the set of parameters θ defined above now only contains λ , or $\theta = \{\lambda\}$. The normalization constant is $Z_\lambda(\Omega) = \sum_{x \in \Omega} x^{-\lambda}$. The derivative with respect to λ of the cross-entropy, $H_{\text{cross}}(f||p(z|\theta, \Omega)) = \lambda \sum_{i=1}^W f_i \log z_i + \log Z_\lambda(\Omega)$, has to be computed, and setting $dH_{\text{cross}}/d\lambda = 0$ yields

$$\sum_{i=1}^W f_i \log z_i = \left(\sum_{i=1}^W z_i^{-\lambda} \right)^{-1} \sum_{i=1}^W z_i^{-\lambda} \log z_i \quad . \quad (9)$$

The solution to this implicit equation, $\lambda = \lambda^*$, can not be written in closed form but can be easily solved numerically. See [32] for the corresponding algorithm and APPENDIX B for a tutorial.

How to determine λ^*

One possibility to find the solution $\lambda = \lambda^*$ from the implicit equation Eq. (9), is to iteratively refine approximate solutions. For this, select $M+1$ values λ from the interval $[\lambda_{\min}, \lambda_{\max}]$, where M is a finite fixed number, say $M = 100$. Those values may be chosen to be given by the expression

$$\lambda_r(m) = \underline{\lambda}_r + \frac{m}{M} (\bar{\lambda}_r - \underline{\lambda}_r) \quad , \quad (10)$$

for $m = 0, \dots, M$. The parameters $\underline{\lambda}_r$ and $\bar{\lambda}_r$ are defined in the following way: First define $\underline{\lambda}_1 = \lambda_{\min}$, and $\bar{\lambda}_1 = \lambda_{\max}$, where λ_{\max} and λ_{\min} are parameters of the algorithm. Then define $\delta\lambda_1 = \Delta\lambda/M$ with $\Delta\lambda = \lambda_{\max} - \lambda_{\min}$. If $\lambda_1(m_1^*)$ is the optimal solution of Eq. (9)




FIG. 3. Comparison of the three power-law exponent estimators, LS, ML_{CSN} , and ML^* . For 400 values of λ in the range between 0 and 4, we sample $N = 10,000$ events from $\Omega = \{1, \dots, 1,000\}$, from a power-law probability distribution $p(x|\lambda, \Omega) \propto x^{-\lambda}$. The estimated exponents λ_{est} for the estimators LS (red), the ML_{CSN} (green, $\lambda_{\text{est}} = \hat{\lambda}$), and the new ML^* (black, $\lambda_{\text{est}} = \lambda^*$), are plotted against the true value of the exponent λ of the probability distribution samples are drawn from. Clearly, below $\lambda \sim 1.5$ the ML_{CSN} estimator no longer works reliably. ML_{CSN} and ML^* work equally well in a range of $1.5 < \lambda < 3.5$. Outside this range ML^* performs consistently better than the other methods. The inset shows the mean-square error σ^2 of the estimated exponents. The LS-estimator has a much higher σ^2 over the entire region, than the ML^* -estimator. The blue dot represents the ML^* estimate for the Zipf exponent of C. Dickens' "A tale of two cities". Clearly, this exponent could never reliably be obtained from the rank ordered distribution using ML_{CSN} , whereas ML^* works fine even for values of $\lambda \sim 0$.

for some m_1^* , then we can choose $\lambda_2 = \lambda_1(m_1^*) - \delta\lambda_1$, and $\bar{\lambda}_2 = \lambda_1(m_1^*) + \delta\lambda_1$ and $\delta\lambda_2 = 2\delta\lambda_1/M$. One then continues by iterating r times until $\delta\lambda_r < \varepsilon$, where ε is the desired accuracy of the estimate of λ^* . As a consequence, the value m_r^* , for which $|\lambda^* - \lambda_r(m_r^*)| < \varepsilon$ holds, optimally estimates λ^* in the r 'th iteration with an error smaller than ε . Note that ε is the error of the ML^* -estimator with respect to the exact value of the predictor λ^* , and is not the error of λ^* with respect to the (typically unknown) value of the exponent λ of the sampling distribution.

Controlling the fit region over which the power-law should be obtained therefore becomes a matter of restricting the sample space to a convenient $\Omega' \subset \Omega$. This can be used for dynamically controlling low-frequency cut-offs. These cut-offs are set to exclude states for which,

$$p(z_i|\lambda, \Omega)N < k_{\min} \quad , \quad (11)$$

where k_{\min} is the minimal number of times that any state- i is represented in the data set. This means that we

re-estimate λ on $\Omega' \subset \Omega$ with

$$\Omega' = \{z_i \in \Omega | p(z_i|\lambda, \Omega)N \geq k_{\min}\} \quad . \quad (12)$$

We see in Eq. (9) that iteratively adapting Ω to subsets Ω' , and then re-evaluating λ , requires to solve,

$$\sum_{i \in I(\Omega')} f'_i \log(z_i) = \left(\sum_{i \in I(\Omega')} z_i^{-\lambda} \right)^{-1} \sum_{i \in I(\Omega')} z_i^{-\lambda} \log z_i \quad , \quad (13)$$

where $N' = \sum_{i \in I(\Omega')} k_i$ is the restricted sample-size and $f'_i = k_i/N'$ are the relative frequencies re-normalized for Ω' . $I(\Omega') = \{i | z_i \in \Omega'\}$ is the index-set of Ω' .

Iterating this procedure either leads to a fixed point or to a limit cycle between two low-frequency cut-offs with two slightly different estimates for λ^* . These two possibilities need to be considered in order to implement an efficient stopping criterion for the iterative search of the desired low-frequency cut-off in the data. The algorithm therefore consists of two nested iterations. The "outer iteration" searches for the low-frequency cut-off, the "inner iteration" solves the implicit equation for the power-law exponent. The matlab code for the algorithm is found in [32], see APPENDIX B for a tutorial.

TESTING THE NEW ESTIMATOR WITH NUMERICAL EXPERIMENTS AND KNOWN DATA SETS

To test the proposed algorithm implementing the estimator ML^* , we first perform numerical experiments and then test its performance on a number of well known data sets.

Testing with numerical experiments

For 400 different values of λ , ranging from 0 to 4, we sample $N = 10,000$ data points $x \in \Omega = \{1, \dots, W\}$, with $W = 1000$ states, with probabilities $p(x|\lambda, \Omega) \propto x^{-\lambda}$. We fit the data in three ways, using (i) least square fits (LS), (ii) the CSN algorithm ML_{CSN} providing estimates $\hat{\lambda}$, and (iii) the implicit ML^* method providing estimates λ^* . In Fig. 3 we show these estimates for the power exponents, as a function of the true values of λ . The LS, ML_{CSN} , ML^* estimators are shown as the red, green, and black curves respectively. Obviously ML^* and ML_{CSN} work equally well for power-law exponents λ with values $1.5 < \lambda < 3.5$. In this range the three approaches coincide. However, note that in the same region the mean square error[?] σ^2 for the LS method is much larger than for ML^* and ML_{CSN} . Outside this range the assumptions and approximations used for ML_{CSN} start to lose their validity and both LS and ML^* estimates outperform the

	exp.	CSN ₁	CSN ₂	ML*	KS _{CSN}	KS*
blackouts	λ	2.3	2.27	2.25	0.061	0.031
surnames	α	2.5	2.49	2.66	0.041	0.019
int. wars	λ	1.7	1.73	1.83	0.078	0.076
city pop.	λ	2.37	2.36	2.31	0.019	0.016
quake int.	λ	1.64	1.64	1.88	0.092	0.085
relig. fol.	λ	1.8	1.79	1.61	0.091	0.095
citations	λ	3.16	3.16	3.10	0.010	0.018
words	α	1.95	1.95	1.99	0.009	0.015
wealth	λ	2.3	2.34	2.30	0.063	0.066
papers	λ	4.3	4.32	3.89	0.079	0.082
sol. flares	λ	1.79	1.79	1.81	0.009	0.021
terr. attacks	λ	2.4	2.37	2.36	0.018	0.017
websites	λ	2.336	2.12	1.72	0.025	0.056
forest fires	λ	2.2	2.16	2.46	0.036	0.034
Dickens novel	λ	-	-	1.04	-	0.017

TABLE I. Comparison of the estimators ML* and ML_{CSN} on empirical data sets that were used in [24]. These include the frequency of surnames, intensity of wars, populations of cities, earthquake intensity, numbers of religious followers, citations of scientific papers, counts of words, wealth of the Forbes 500 firms, numbers of papers authored, solar flare intensity, terrorist attack severity, numbers of links to websites, and forest fire sizes. We added the word frequencies in the novel “A tale of two cities” (C. Dickens). The second column states if α or λ were estimated. The exponents reported in [24] are found in column CSN₁, those reproduced by us applying their algorithm to data [24, 27–30] is shown in column CSN₂. The latter correspond well with the new ML* algorithm. For values $\lambda < 1.5$, CSN can not be used. We list the corresponding values for Kolmogorov-Smirnov test for the two estimators, KS_{CSN} and KS*.

ML_{CSN} estimates. The inset also shows that ML* consistently estimates λ much better than the LS estimator (two orders of magnitude better in terms of σ^2) for the entire range of λ . The blue dot in Fig. 3 represents the ML* estimate for the Zipf exponent of C. Dickens’ ‘A tale of two cities’. Clearly, this small exponent could never be obtained by ML_{CSN}, see also Tab. I.

Testing with empirical data sets

We finally compare the new estimator ML* on several empirical data sets that were used for demonstration in [24]. In Tab. I we collect the results. The second column states if λ or α were estimated. Column CSN₁ presents the value of the estimator ML_{CSN} as presented in [24]. Column CSN₂ contains the values of the same estimator using the data from [24] and using the algorithm provided by [31][?]. The results for the ML* estimator agrees well with those of ML_{CSN} in the range where the latter works well. To demonstrate how ML* works perfectly outside of the comfort zone of ML_{CSN} (for $\lambda < 1.5$), we add the result of the rank distribution of word counts in the novel “A tale of two cities” (Charles Dickens, 1859),

which shows an exponent of $\lambda \sim 1.035$. This exponent can be fitted directly from the data using the proposed ML* algorithm, while ML_{CSN} can not access this range, at least not without the detour of first producing a histogram from the data and then fitting the tail of the frequency distribution. The values for the corresponding Kolmogorov-Smirnov tests (see e.g. [24]) for the two estimates, KS_{CSN} and KS*, are similar for most cases.

CONCLUSIONS

We discuss the generic problem of estimating power-law exponents from data sets. We list a series of questions that must be clarified before estimates can be performed. We present these questions in form of a decision tree that shows how the answers to those questions lead to different strategies for estimating power-law exponents.

To follow this decision tree can be seen as a recipe for fitting power exponents from empirical data. The corresponding algorithms were presented and can be downloaded as matlab code. The two algorithms we provide are based on a very general ML estimator that maximizes an appropriately defined cross entropy. The method can be seen as a straight forward generalization of the idea developed in [24]. The two estimators (one for binned histograms and ML* for raw data sets) allow us to estimate power-law exponents in a much wider range than was previously possible. In particular, exponents lower than $\lambda < 1.5$ can now be reliably obtained.

ACKNOWLEDGMENTS

This work was supported in part by the Austrian Science Foundation FWF under grant P29252. B.L. is grateful for the support by the China Scholarship Council, file-number 201306230096.

* stefan.turner@meduniwien.ac.at

- [1] M.E.J. Newman, *Power-laws, Pareto distributions and Zipf’s law*, Contemporary physics 2005; **46** 323–51.
- [2] M. Mitzenmacher, *A Brief History of Generative Models for Power-Law and Lognormal Distributions*, Internet Mathematics 2004; **1** 226–51.
- [3] L.P. Kadanoff, et al., *Static Phenomena Near Critical Points: Theory and Experiment*, Rev. Mod. Phys. 1967; **39** 395–413.
- [4] D. Sornette, *Critical Phenomena in Natural Sciences*, Springer, Berlin, 2006.
- [5] P. Bak, C. Tang, and K. Wiesenfeld, *Self-Organized Criticality: An Explanation of 1/f Noise*, Phys. Rev. Lett. 1987; **59** 381–84.
- [6] H.A. Simon, *On a class of skew distribution functions*, Biometrika 1955; **42** 425–40.

- [7] A.Réka, and A.L. Barabási, *Statistical mechanics of complex networks*, Rev. Mod. Phys. 2002; **74** 47–97.
- [8] A.L. Barabási, and A. Réka, *Emergence of scaling in random networks*, Science 1999; **286** 509–12.
- [9] G.U. Yule, *A Mathematical Theory of Evolution, based on the Conclusions of Dr. J. C. Willis, F.R.S*, Phil. Trans. Royal Soc. B 1925; **213** 21–87.
- [10] H. Takayasu, A.-H. Sato, and M. Takayasu, *Stable Infinite Variance Fluctuations in Randomly Amplified Langevin Systems*, Phys. Rev. Lett. 1997; **79** 966–67.
- [11] C. Tsallis, *Introduction to nonextensive statistical mechanics*, Springer, New York, 2009.
- [12] R. Hanel, S. Thurner, S, and M. Gell-Mann, *How multiplicity of random processes determines entropy: derivation of the maximum entropy principle for complex systems*, Proc. Nat. Acad. Sci. USA 2014; **111** 6905–10.
- [13] B. Corominas-Murtra, R. Hanel, and S. Thurner, *Understanding scaling through history-dependent processes with collapsing sample space*, Proc. Nat. Acad. Sci. USA 2015; **112**, 5348–53.
- [14] B. Gutenberg, and C.F. Richter, *Frequency of earthquakes in California*, Bull. Seismol. Soc. Amer. 1944; **34** 185–88.
- [15] K. Christensen, L. Danon, T. Scanlon, and P. Bak, *Unified scaling law for earthquakes* Proc. Nat. Acad. Sci. USA 2002; **99** 2509–13.
- [16] F. Auerbach, *Das Gesetz der Bevölkerungskonzentration*, Petermanns Geographische Mitteilungen 1913; **59** 74–76.
- [17] X. Gabaix, *Zipf’s Law for Cities: An Explanation*, Quart. J. Econ. 1999; **114** 739–67.
- [18] C.A. Shaffer, *Spatial foraging in free ranging bearded sakis: Traveling salesmen or Lévy walkers?*, Amer. J. Primatology 2014; **76** 472–84.
- [19] M.E.J. Newman, and R.G. Palmer, *Modeling extinction*, Oxford University Press, 2003.
- [20] G.K. Zipf, *Human Behavior and the Principle of Least Effort*, Addison-Wesley, Cambridge, Massachusetts, 1949.
- [21] S.J. Press, *Subjective and Objective Bayesian Statistics: Principles, Models, and Applications*, Wiley Series in Probability and Statistics, 2010.
- [22] J.O. Berger, *Statistical decision theory and Bayesian Analysis*, Springer, New York, 1985.
- [23] R.A. Fisher, *On an absolute criterion for fitting frequency curves*, Messenger of Mathematics 1912; **41** 155–60.
- [24] A. Clauset, C.R. Shalizi, and M.E.J. Newman, *Power-Law Distributions in Empirical Data*, SIAM Review 2009; **51** 661–703.
- [25] Y. Virkar, and A. Clauset, *Power-law distributions in binned empirical data*, Annals of Applied Statistics 2014; **8** 89–119.
- [26] A. Deluca, and A. Corral, *Fitting and goodness-of-fit test of non-truncated and truncated power-law distributions* Acta Geophysica 2013; **61** 1351–94
- [27] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener, *Graph structure in the web*, Computer networks 2000; **33** 309–20.
- [28] D.C. Roberts, and D.L. Turcotte, *Fractality and self-organized criticality of wars*, Fractals 1998; **6** 351–57.
- [29] S. Redner, *How popular is your paper? An empirical study of the citation distribution*, EPJ B 1998; **4** 131–34.
- [30] A. Clauset, M. Young, and K.S. Gleditsch, *On the frequency of severe terrorist events*, Journal of Conflict Resolution 2007; **51** 58–87.
- [31] <http://tuvalu.santafe.edu/~aaronc/powerlaws/>
- [32] http://www.complex-systems.meduniwien.ac.at/SI2016/r_plfit.m
http://www.complex-systems.meduniwien.ac.at/SI2016/r_plhistfit.m
- [33] H.S. Heaps, *Information Retrieval: Computational and Theoretical Aspects*, Academic Press, 1978.
- [34] G. Herdan, *Type-token mathematics*, Gravenhage, Mouton & Co, 1960.

APPENDIX A: SAMPLING FROM CONTINUOUS SAMPLE SPACES

If events x are drawn from a continuous sample space $\Omega = [x_{\min}, x_{\max}]$, for instance the magnitude of earthquakes, then the ‘natural order’ of possible events is simply given by the magnitude x of the observation. Events x are drawn from a continuous power-law distribution $p(z|\lambda, \Omega) = x^{-\lambda}/Z$, with $Z = Z_\lambda([x_{\min}, x_{\max}])$ (compare Eq. (3) first line).

To work with well defined probabilities we have to bin the data first. Probabilities to observe events within a particular bin depend on the margins of the W bins $b = (b_0, b_1, \dots, b_W)$, with $b_0 = x_{\min}$ and $b_W = x_{\max}$. The histogram $k = (k_1, \dots, k_W)$ counts the number k_i of events x falling into the bin $b_i > x \geq b_{i-1}$, and the probability of observing x in the i ’th bin is given by

$$p(i|\lambda, x) = \frac{b_i^{1-\lambda} - b_{i-1}^{1-\lambda}}{x_{\max}^{1-\lambda} - x_{\min}^{1-\lambda}}. \quad (14)$$

Binning events sampled from a continuous distribution may have practical reasons. For instance data may be collected from measurements with different physical resolution levels, so that binning should be performed at the lowest resolution of data points included in the collection of samples. We will not discuss the ML estimator for binned data in detail but only remark that for given bin margins b it is sufficient to insert $p(i|\lambda, x)$ of Eq (14) into Eq. (7) with $\theta = \{\lambda\}$, to derive the appropriate ML condition for binned data. An algorithm for binned data `r_plhistfit`, where we assume the bin margins b_i to be given, is found in [32].

We point out that if margins for binning have not been specified prior to the experiments, then specifying the optimal margins for binning the data becomes a parameter estimation problem in itself, i.e. the optimal margins b_i have to be estimated from the data as well. One major source of uncertainty in the estimates of λ from binned data is related to the uncertainty in choosing the upper and lower bounds x_{\min} and x_{\max} of the data, i.e. specifying the bounds of the underlying continuous sample space.

Binning becomes irrelevant for clean continuous data for the following reason. Suppose we fix the sample space $[x_{\min}, x_{\max}]$ and cut this domain into M bins of width $\Delta = (x_{\max} - x_{\min})/M$. Since the data $x = \{x_1, \dots, x_N\}$ is drawn from a continuous sample space, the chance for two observations x_m and x_n to be exactly equal becomes zero for $m \neq n$, if M has been chosen sufficiently large. Then each bin almost certainly contains either one sample x_n or none. The probability of observing x then is asymptotically (as Δ approaches zero) given by

$$P(x|\lambda) = \Delta^N \prod_{n=1}^N \left(\frac{x_n^{-\lambda}}{Z_\lambda(x_{\min}, x_{\max})} \right). \quad (15)$$

The parameter estimation problem of finding the optimal λ is equivalent to maximizing $P(x|\lambda)$ (or equivalently $\log P(x|\lambda)$) with respect to λ . In this maximization problem Δ becomes irrelevant and only the choice of x_{\min} and x_{\max} and the data x remains relevant for the estimate. As a consequence, one obtains an equation

$$\sum_{i=1}^W f_i \log z_i = \frac{d}{d\lambda} \log Z_\lambda, \quad (16)$$

for the ML estimate of the exponent λ over continuous sample spaces. Equation (9) and Eq. (16) differ only in Z_λ . In Eq. (9) the normalization constant of discrete samples spaces gets used while in Eq. (16) Z_λ is the normalization constant for a continuous sample space. Switching between continuous and discrete sample spaces therefore is simply a matter of choosing the one or the other normalization constant in the algorithm.

Whether data should be assumed to be sampled from continuous or discrete sample spaces is not always totally clear. Many measurements have an intrinsic resolution and implicitly bin the data. For instance if real numbers sampled in an experiment are given only with a three digit precision, such as $x_n = 0.123$ and we know that $0.001 = x_{\min}$ and $x_{\max} = 5$ then we better treat the data as discrete data on $\Omega_d = \{0.001, 0.002, \dots, 4.998, 4.999, 5\}$ if we have sufficiently many samples for the histogram over Ω_d not to be flat. A primitive test to see whether one should regard data as sampled from a continuous sample space or not is to make a histogram over the unique values of the recorded data. If each distinct value appears only once in the data (i.e. if the histogram over the unique data-points is flat) then one should treat the sample-space as continuous.

While for the discrete case we need not estimate x_{\min} and x_{\max} this remains necessary for the continuous case. The method of cutting the $[x_{\min}, x_{\max}]$ into segments of length Δ and then taking Δ to zero explains why typically the *primitive* estimates, $x_{\min} = \min\{x_n | n = 1, \dots, N\}$ and $x_{\max} = \max\{x_n | n = 1, \dots, N\}$, provides fairly good results. Alternatively, strategies such as suggested in [24] could be used to optimize the choices for x_{\min} and x_{\max} . However, this procedure can not be directly derived from Bayesian arguments. Neither will we discuss this approach in this paper nor implement such an option in `r_plfit`.

However, Bayesian estimates of x_{\min} and x_{\max} exist. Although we will not discuss those estimators in detail here we will eventually implement them in `r_plfit` to replace the primitive estimates. The idea of constructing such estimators is the following. For instance, one asks how likely can the maximal value $\max(x)$ of the sampled data $x = (x_1, \dots, x_N)$ be found to be larger than some value y . By deriving $P(\max(x) > y|\lambda, [x_{\min}, x_{\max}])$ and $P(\min(x) < y|\lambda, [x_{\min}, x_{\max}])$, as a consequence, it becomes possible to derive Bayesian estimators for x_{\min}

and x_{\max} .

APPENDIX B: USING R.PLFIT

The matlab function

```
function out = r_plfit(data,varargin)
```

implements the algorithm discussed in the main paper. The function returns a struct `out` that contains information about the data, the data range, but most and for all `out.exponent` returns the estimated exponent of the power-law. Whether the exponent `out.exponent` is the exponent λ of the sample distribution or the exponent α of the frequency distribution of the data depends on how `function out = r_plfit(data,varargin)` gets used as explained below. In the code the sample space Ω is equivalent to a vector $z = [z_1, \dots, z_W]$ containing W distinct event magnitudes z_i , $i = 1, \dots, W$.

The variable `data` can be used to import data while a variable number of arguments can be set by `varargin` to tell the algorithm which type of data it should handle and to control the range of the data. By default the only argument that has to be set is `data`. `r_plfit` filters data from data points `data <= 0`, `NaN`, `Inf`. The data passed on to `data` can be

- a vector of observations `data` $\equiv x = [x_1, \dots, x_N]$ (default)
- a histogram `data` $\equiv k = [k_1, \dots, k_W]$ of recorded event types $i = 1, \dots, W$

`out = r_plfit(data,varargin)` can be used in three basic modes

- `out = r_plfit(x)` returns the estimated exponent λ of the **probability distribution** given the observation x (default)
- `out = r_plfit(k,'hist')` returns the estimated exponent λ of the **probability distribution** given the histogram of observations k
- `out = r_plfit(k)` returns the estimated exponent α of the **frequency distribution** given the histogram of observations k

The third mode `out = r_plfit(k)` is in fact identical to the first mode `out = r_plfit(x)`, only that passing a histogram as sample data to the algorithm is identical to asking how many of the W states i have been observed n times. But this is exactly the frequency distribution of the process, which possesses a tail with exponent $\alpha = 1 + 1/\lambda$. Depending on the mode `r_plfit` returns the exponent λ or α in `out.exponent`

Fitting with observations x : If we run `out = r_plfit(x)` without further options `r_plfit` assumes by default that the data x consists of natural numbers, and that the process samples have been sampled from the sample space $\Omega = \{\min(x), \min(x) + 1, \dots, \max(x) - 1, \max(x)\}$, i.e. $\min(x) \leq z_i = i \leq \max(x)$. If this is not the case one can either specify the data range using all W unique values $z = [z_1, \dots, z_W]$ occurring in the data x by using the option `out = r_plfit(x,'urange')`. In order to define a fit range maximal and minimal data values taken into account can be set by `out = r_plfit(x,'urange','rangemin',minval, ... 'rangemax',maxval)` such that `r_plfit` only takes into account data in the range $\minval \leq z \leq \maxval$. To control the data range individually use `out = r_plfit(x,'range',z)`. If the data has been sampled from a continuous sample space, and the histogram over the unique data is flat, i.e. each value in the data only appears once (more or less), then one can tell `r_plfit` that the data is sampled from a continuous sample space by setting the option `'cdat'`, i.e. by running `out = r_plfit(x,'cdat', ...)`. This option tells the algorithm to use the normalization constant for continuous sample spaces and estimates $x_{\min} = \min(x)$ and $x_{\max} = \max(x)$. Moreover, `'cdat'` implicitly sets the `'urange'` and the `'nolf'` option. `'nolf'` (see below) switches off the search of the algorithm for an optimal low frequency cut-off.

Fitting with histograms k : Using histograms k as input works in exactly the same way as for fitting x if we want to estimate the exponent α of the frequency distribution and use `r_plfit` in the `out = r_plfit(k)` mode. If we use `r_plfit` in the `out = r_plfit(k,'hist')` mode, the algorithm assumes by default that the sample space z is given by $z = [1, 2, \dots, W]$. The option `'urange'` has no effect in this mode and gets ignored if set. Otherwise one can again use the `'range'` property to set the event magnitudes z (the sample space) using `out = r_plfit(k,'hist','range',z)`. The `'minrange'` and `'maxrange'` options work in exactly the same way as before.

Dynamic low frequency cut-off: By default `r_plfit(data)` runs an iterative search for an optimal low frequency cut-off that is set at a range value z_i such that the expected number of samples for z_i equals the variable N_{\min} (default value 1, reset using option `'Nmin'`). This means the algorithm performs a low frequency cut-off for observations x . If however `maxval` is smaller than the predicted cut-off then the low frequency cut-off has no effect. One should note that in the mode `out = r_plfit(k)` the low frequency cut-off mechanism effectively acts as a high frequency cut-off with respect to the data x . One can switch this mechanism off by

setting the option 'nolf' (no low frequency cut-off).

USING R_PLHISTFIT

The 'plot' option, `out = r_plfit(data,..., 'plot')`, can be used for visualization. `r_plfit` plots the fit over the data in double logarithmic coordinates (loglog plot). Using the option 'figure' behaves like 'plot' but explicitly opens a new figure. 'exp_min' can be used to specify the minimal search value for the exponents (default is 0) and 'exp_max' to set the maximal search value (default is 5). 'eps' can be used to set the precision of the implicit algorithm (default $1e-5$). Several other options exist to control the performance of the algorithm, which all can be listed by using `r_plfit('help')` in the command line, which prints a brief manual on the usage of `r_plfit` and available options.

If one works with binned data, e.g. histogram data counting the number of events falling into exponentially scaled bins (log-binning), then `r_plhistfit` needs to be used instead of `r_plfit`. The function `function out = r_plhistfit(data,varargin)` like `r_plfit`, by default, uses only data as input and other variables can be set optionally. `data` is always a histogram k that is a vector $k = [k_1, \dots, k_W]$. Bins can be specified by giving bin margins $b = [b_0, b_1, \dots, b_W]$ such that events counted in k_i had a magnitude x such that $b_{i-1} \leq x < b_i$. Usage, `r_plhistfit(k, 'margins', b)`. By default `r_plhistfit` assumes that $b_i = i+1/2$. Other options work similar to the ones available for `r_plfit` and can be reviewed by typing `r_plhistfit('help')` in the matlab command line.